

Framework for DoD's Optical Networking Security Standards Effort

Abstract

This document presents a framework for developing security standards for optical networks. It is based on an assessment of the emerging capabilities of optical networks and their role in supporting the Global Information Grid (GIG). This document:

- Discusses DoD's requirements for hardening the infrastructure of the GIG's networks,
- Presents a snapshot of current optical networking standards with high-level descriptions of the main protocols supporting optical networking,
- Compares and contrasts the standards organizations involved in developing optical networking standards as potential venues for incorporating DoD's optical networking security standards, and
- Concludes with a roadmap for the optical network security standards work.

1. Introduction

1.1 Background

Optical fiber communications and emerging Dense Wavelength Division Multiplexing (DWDM) technology offers an abundance of bandwidth and additional flexibility for configuring this bandwidth across mesh topologies containing a combination of electro-optical and all-optical components. Currently, a single fiber can support up to 16 wavelengths (lambdas), each operating at an OC-48 (2.5 Gbit/s) rate. Because of the development of improved DWDM components, available bandwidth per fiber is expected to double every twelve months and reach 16 Terabit/s within the foreseeable future.

To use this explosive growth of bandwidth effectively, the industry is working on rapid provisioning systems. Instead of long provisioning times, typically two to four weeks but sometimes as much as three to six months, engineers are designing protocols for real-time or near-real-time provisioning of circuits.

Until now, circuits have been provisioned manually across transmission equipment that lacks routing and switching protocols for doing this automatically in real time. Multiple Standards Development Organizations (SDOs), such as the Internet Engineering Task Force (IETF) and the International Telecommunication Union (ITU-T) [should the OIF be included here?], are working on routing and switching protocols for real-time, automatic provisioning of lambdas, wavebands (groups of lambdas), and fibers. The momentum behind these efforts is focused on using the Internet Protocol (IP) suite and label switching based on MPLS as the core protocols to control, manage, and operate these optical networks. When realized, this capability to support huge amounts of bandwidth on demand will be a major component of the GIG.

Unfortunately, but not surprisingly, these standards are being developed with little or no consideration for security mechanisms to protect optical networks, and what little security that has been included falls far short of DoD's requirements.

1.2 Purpose

The purpose of this paper is to present a framework for developing optical network security standards that will support DoD's needs for infrastructure hardening. Development of the necessary security standards, including remedies for deficiencies, is beyond the scope of this paper but is expected to follow from the directions set in this paper.

Developing security standards for optical networks is challenging[I think we can easily come up with a technical approach. The overwhelming and confusing part is the politics of finding commercial applicability for meeting DoD's requirements and convincing the SDO's to accept our approach. Suggest striking the overwhelming and confusing text.] due to multiple factors that affect optical networking security standards. First, optical networking is a rapidly moving target. The technologies involved and associated standards are currently being specified as competing products are developed simultaneously. Technical solutions for optical networking often change. Second, network security is a complicated subject, and it is also undergoing changes, even though not as fast as optical networks. Third, multiple SDOs are involved in this effort, and each has different priorities, working methods, and points of view. And finally, many of the standards being considered for optical networks have evolved from existing protocols, and these often contain security mechanisms that either fall short of meeting DoD's requirements or are not uniformly and consistently implemented throughout the full optical networking protocol suite.

In addition to the standards activities addressed by this framework, NSA has developed a Common Criteria Protection Profile (PP) identifying the security requirements and level of assurance needed for a Hardened Switch/Router. The PP specifies commercial-grade security mechanisms at the Evaluated Assurance Level of 3+. The Hardened Switch/Router PP is currently in draft and applies to ATM, IP, and optical switches and routers. The PP will allow commercial testing and certification of hardened routers and switches in a laboratory certified by the National Information Assurance Partnership (NIAP). NSA plans to issue RFPs, starting in FY-02, for the development of switches and routers meeting the NSA's Hardened Switch/Router PP and implementing the applicable security mechanisms identified in the networking standards documents.

This paper attempts to present a common understanding of the problem space and a starting point toward development of the necessary optical networking security standards. The security standards developed through this effort, along with the security requirements identified in the Hardened Switch/Router PP, will allow DoD to partner with industry to develop hardened optical switching components to support DoD's high availability networking needs.

2. DoD's Requirements

Although various requirements for the security of optical networking are essential and unique to DoD, the window of opportunity to influence emerging optical networking standards is rapidly closing. Because time is of the essence, infrastructure protection is the most urgent and utmost requirement. As mentioned above, optical networking technologies offer many potential capabilities required by the GIG and are expected to be the backbone not only for DoD but also for major nationwide backbone networks. Infrastructure protection is aimed at increasing the availability and robustness of these networks by isolating DoD users from other users of the common optical infrastructure and by protecting the control and management components from attacks by individual intruders, by other users of the network, and by organized and sophisticated hostile entities. Isolation of DoD users includes capabilities such as secure and robust virtual private networking, quality of service, and traffic engineering which are aimed at insuring that DoD networking services can not be degraded by other users of the network. The goals are:

1. To support the availability of optical networking communications to DoD users by preventing denial (or degradation) of service attacks originating from outsiders or from other users of the common optical infrastructure.
2. To help isolate DoD users' networks from external analysis of network facilities and configuration and DoD users' communications from traffic analysis.
3. To protect network infrastructure component control and management information from unauthorized external analysis and modification.
4. To provide support for additional security services consistent with the principle of defense in depth.

Although optical networking comprises an emerging set of technologies, the requirements and mechanisms needed for hardening optical networking components are similar to those needed for other networking technologies, such as ATM and IP.

2.1 Three Planes of Optical Networks

Traffic on optical networks, like most networks, can be decomposed into three basic types of communication: user, control, and management.

- User traffic is simply the information that users are transmitting over the network.
- Control traffic is any information transferred between network components that is necessary for establishing user connections. Control traffic includes name resolution, address resolution, route establishment, and signaling. Name resolution (e.g., the Domain Name System, DNS) retrieves relationships like domain names to network addresses. Address resolution (e.g., Address Resolution Protocol, ARP) maps network addresses between protocols or protocol layers. Route establishment or routing (e.g., RIP, OSPF, IS-IS, or BGP) communicates network topology information and uses this information to determine forwarding paths. Signaling (e.g., CR-LDP or RSVP-TE) is used to establish and tear down paths.
- Management traffic consists of any protocols used to configure network components or to inform other components or network administrators about the status of network components. This also includes protocols for neighbor and service discovery (e.g.,

LMP). Other management protocols include Simple Network Management Protocol (SNMP), Hypertext Transfer Protocol (HTTP), the rlogin and telnet command-line interfaces, and proprietary management protocols.

2.2 Requirements for the User Plane

Protecting the confidentiality and integrity of user information is outside the scope of this framework; however, mechanisms to isolate DoD users' traffic from other users who share the common optical infrastructure falls within the scope of this framework. These mechanisms may include non-cryptographic virtual private networking (VPN), traffic prioritization, and traffic engineering capabilities as well as hardened optical infrastructure components. The VPN, traffic prioritization, and traffic engineering mechanisms are aimed at ensuring that service level agreements can be enforced robustly by service providers, so that other users of the network cannot adversely affect the availability of DoD's leased networking services. The hardened optical infrastructure components are aimed at ensuring that an adversary cannot maliciously modify the VPN, traffic prioritization, and traffic engineering parameters that provide isolation of DoD's networking services from other users of the network. Note that other users of the network may include intruders, adversaries who have leased services on the optical infrastructure, or anyone who can access the optical infrastructure through other networks such as the Internet.

2.3 Requirements for the Control and Network Management Planes

Proper name resolution, address resolution, route establishment, and signaling by the network infrastructure are essential for user traffic to be directed to the intended destination. Name and address resolution must be protected to prevent various spoofing attacks and to hide details about the network configuration. Routing information must be protected to ensure that the path user information takes is not manipulated and to conceal updates to the network configuration. Manipulation may include altering the path that user information traverses (e.g., routing traffic through a listening post), altering the quality of service for the path that the information traverses (e.g., degrading service), or eliminating routes for a particular user or group of users (i.e., denying service). Similarly, signaling (i.e., in MPLS-based optical networks, label distribution) must be protected between optical infrastructure components to ensure user connections are established to the proper destinations and not subsequently altered in unauthorized ways. To protect the network's control protocols from unauthorized modification and traffic analysis, DoD requires the authentication, integrity, and confidentiality of routing, signaling, and addressing information. Two-way authentication enables senders and receivers to verify the destination and source of information, respectively. Integrity protects the control information from alteration, rearrangement, replay, truncation, or other such falsification. Confidentiality of the control information protects against network mapping and traffic analysis.

Protecting network management traffic is essential to ensure that network components are not modified by unauthorized users. If management of a network component is compromised, that component can be configured to perform any function the attacker chooses. This includes modification of the network control information and security parameters used to protect the control and user information. Simply being able to view

configuration information on a network component may give an attacker knowledge of network connectivity, addressing schemes, usage statistics, security policies and parameters, and other potentially sensitive information. This information may reveal network configurations, allow traffic analysis, alert an enemy to changes in the network, or facilitate a complete break-in. In this framework, network management includes local or remote management from a network management workstation (e.g., SNMP based HP Openview) via a serial port, LAN, or other out-of-band network connection.

Using commercial networks to carry DoD's traffic may increase the information assurance (IA) requirements on the commercial network's infrastructure. Control and management information, as listed in the examples above, contains critical and sensitive information and therefore requires integrity and confidentiality protection.

2.4 DoD's Requirements for a Complete Set of Security Services

In the previous section, two-way authentication, integrity, and confidentiality were identified as required security services for optical control and management protocols. Certain control and management protocols for optical networking are specified with a security option like a message integrity check based on a particular keyed hash function. For several reasons, such features are often problematical for DoD's purposes. For example:

1. Other required services like replay detection and confidentiality may not be available.
2. Certain attacks, like truncation of the data stream, may not be addressed.
3. Authentication and key management, which are far more difficult problems, may not be specified at all.
4. The user may have no way to choose which services are needed for a given deployment of a protocol.
5. A fixed crypto algorithm may be specified, so the user has no way to choose alternatives, upgrade the system as needed, or deploy user-defined algorithms.
6. Procedures for providing security services in proxy devices or for nesting security services across different network segments or administrative domains may be lacking.
7. Methods for specifying and enforcing security policy may be missing.

A list of requirements is shown in Annex A. This list is intended to clarify the capabilities, flexibility, and supporting security services DoD needs to deploy switches and routers with appropriately hardened network infrastructure protocols.

3. Optical Networking Protocols

3.1 IETF's View

The IETF usually views itself as the entity that specifies protocols but does not try to determine how these protocols are used. Therefore, the architecture [draft-ietf-ccamp-gmpls-architecture-00.txt] and framework documents [draft-many-ccamp-gmpls-framework-00.txt] do not contain any reference diagrams for the optical network

architecture, so it is difficult to identify network elements and interface types among them. In addition, the IETF's view on optical networking is still emerging.

The architecture document does enumerate the protocols that are needed to support GMPLS-based optical networks. In addition to GMPLS, which is based on RSVP-TE and CR-LDP, OSPF-TE and IS-IS-TE are used for intra-domain routing. Link Management Protocol (LMP) is used to manage the control channel and a large number of links between a pair of nodes. It is expected that BGP will be extended for inter-domain routing.

The architecture document mentions briefly, in Section 3.3, three different network models: an overlay model, an augmented model, and a peer (integrated) model. An ATM network is an example of the overlay model, in which network nodes do not share topology information with end systems, and networking protocols are different from access protocols. A peer model, in which network nodes share topology information with end systems and the UNI and NNI protocols are same, is the opposite of an overlay model.

However, the same document indicates the IETF's indifference on this subject. Section 7 supports a peer-only model or both. Regardless, IETF has a culture against service providers, and especially against traditional network operators such as Regional Bell Operating Companies.

3.2 OIF's View

The OIF aligns with IETF in terms of the protocols used for optical networks. The UNI 1.0 specification is based on RSVP-TE and CR-LDP, and it includes service discovery and neighbor discovery, which are based on LMP. Figure 1 shows the OIF's view of the optical networking architecture.

Although the OIF is IP-centric in its selection of control and management protocols, it differs significantly from IETF in its networking model. As indicated in its UNI 1.0 work, it clearly supports an overlay model. OIF is now conducting a principal ballot the UNI 1.0 specification and is investigating whether a NNI specification should be developed. The decision to develop an NNI specification is expected in November 2001.

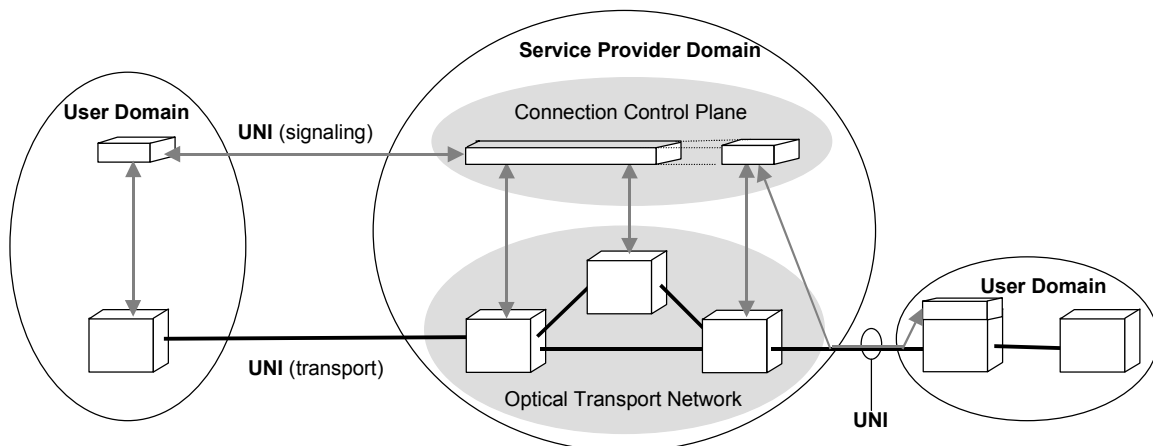


Figure 1: OIF's Reference Architecture.

3.3 ITU-T's view

The ITU-T's view on optical network architecture is shown in Figures 2. They are based on G.ASTN Automatic Switched Transport Network, which is being developed by SG 13.

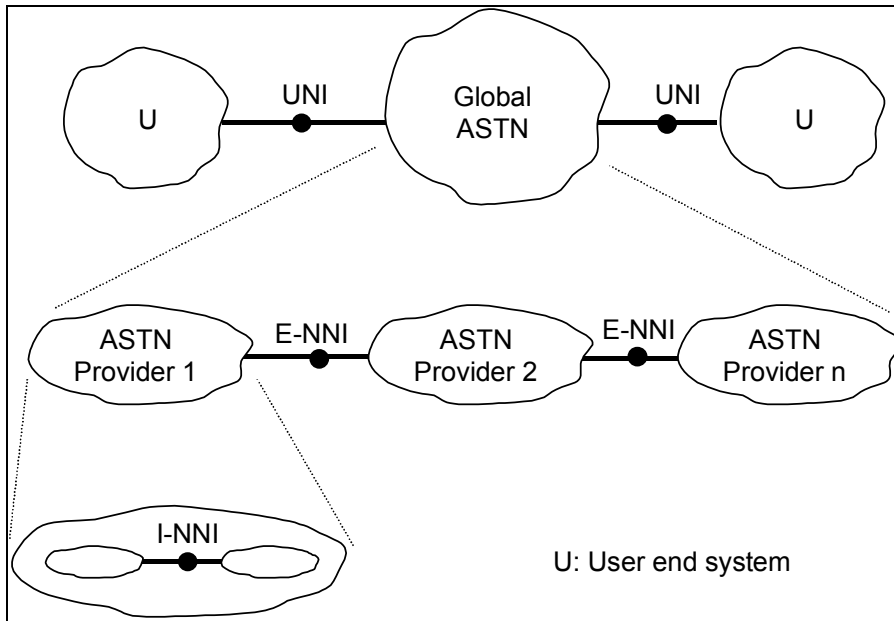


Figure 2. ASTN Architecture.

As shown in Figure 2, ITU's view is similar to that of OIF. As ITU has been throughout the years, it identifies separate protocols for UNI and NNI. It appears that NNI may be further divided into Internal NNI and External NNI.

Figure 3 shows work areas of ASTN. Development of each area should be monitored and draft recommendations should be reviewed for analysis of deficiencies.

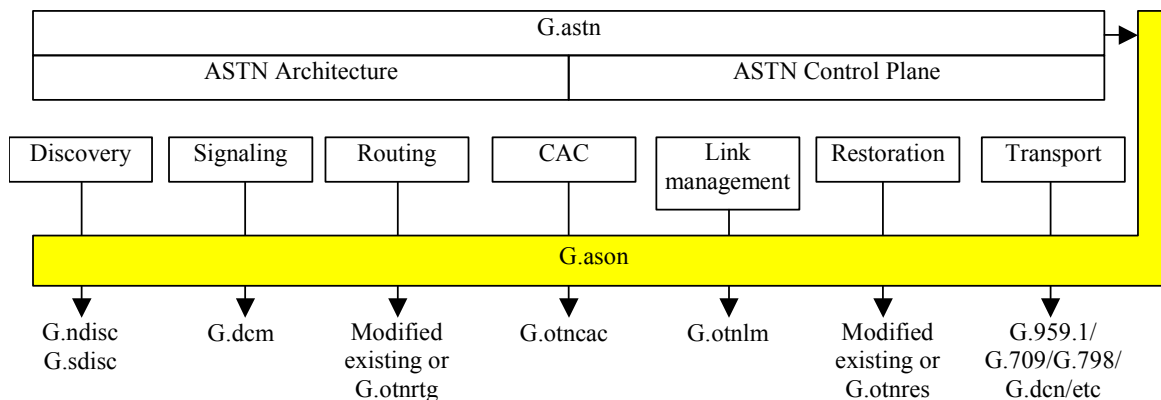


Figure 3. ASTN Work Areas

4. Assessment of Standards Organizations

Optical networking standards are being developed by multiple organizations. Among them, the IETF, OIF, and ITU-T are currently the most important players. (ANSI technical subcommittee T1X1 also works on optical networking standards but serves mainly as a feeder organization to ITU-T from the U.S.)

Although these three organizations are coordinating their optical network standards work, there exist significant overlaps among their work.

4.1 IETF

Optical standards

IETF working groups in the Sub-IP Area such as Common Control and Measurement Plane (ccamp), IP over Optical (ipo), and Multiprotocol Label Switching (mpls) are involved in optical networking standards.

- The ccamp Working Group was created to provide a common approach for controlling sub-IP technologies with the IP protocol suite. This working group is responsible for developing Generalized MPLS, where “generalized” refers to controlling not only packet switching but also circuit technologies such as TDM, wavelength, waveband, and fiber switching.
- The ipo Working Group is responsible for identifying requirements specific to the optical layer so that they can be used by ccamp to develop GMPLS.
- The mpls Working Group initially developed MPLS and is now working on operations and maintenance functions such as fast restoration of MPLS LSPs (label switched paths).

Security standards

The IETF’s Security Area has been producing standards for about 12 years. Various working groups have developed a large number of security standards at different protocol layers. For example, PEM, S/MIME, and OpenPGP define secure file or message formats. Kerberos and GSS-API are application-layer security protocols. TLS and SSH

define security at a session layer. IPsec is a general network-layer security system. DNSSec is an infrastructure security protocol. PKIX and SPKM address long-term key management. However, many protocols (e.g., SNMP, PPP, RSVP, LDP) have security specifications that were written outside of the Security Area, and no guidance is given as to what protocols should be used where. Consequently, development of security standards is fragmented, and there is no security architecture specifying how to secure the network infrastructure. DoD may need to fill this gap with its own architecture for hardening the infrastructure protocols used in optical networking.

The IPsec protocol provides a comprehensive approach to security that can be adapted or profiled to satisfy most, if not all, of DoD's requirements. On the other hand, several of the important protocols for optical networking (e.g., RSVP and LDP) address security by defining only a message authentication code based on a shared secret key and a fixed hash function. The shortcomings of this approach are enumerated in detail in Section 2, above, but some vendors have justified the latter approach by citing the perceived complexity and expense of implementing IPsec. Because of this, the most problematical aspect of IPsec, the key management protocol IKE, is currently being redesigned, but it is too early to draw any conclusions from this effort.

Assessments of IETF

The IETF is arguably the most important communications standards organization. Because the communication world is seemingly moving toward either all-IP or IP-centric designs, and the IETF is the organization responsible for IP and Internet standards, the IETF is viewed as the most influential communications standards organization.

However, the IETF process for standards development is more difficult to predict than those of the ITU-T and leading industry forums such as the ATM Forum. Unlike ITU-T's consensus-driven process, or the ATM Forum's or OIF's voting process, key decisions in the IETF are made by a few "wise men," called Area Directors (ADs). The IETF Chair and ADs make up the Internet Engineering Steering Group (IESG), which makes decisions such as what work items to accept, which specifications progress to the next level, etc. This is not an open, transparent process. For organizations without a foothold in the IESG, it is often difficult to tell what is happening and to know what to expect.

The IETF has a strong culture that favors decentralization. Leaders of IETF tend to dislike intelligence in the network and want to promote intelligence in the user devices or host machines. For this reason, service providers are not favorably received in IETF and are not as influential as they are in the ITU-T or industry forums.

Target opportunities

At the last IETF meeting held in July 2001, the ccamp Working Group decided to start a design team for security for GMPLS. This may be the beginning of a *bona fide* attempt to address deficiencies in optical networking security. DoD can ill-afford to miss this opportunity to influence the optical networking security standard work from the beginning.

Possible near term actions by DoD include:

- Complete the analysis of security mechanisms in existing networking protocols to baseline activities for a draft document on comprehensive security standards for infrastructure protection.
- Participate on the ccamp security design team.
- Write an Internet Draft on security requirements including needs for comprehensive security standards for infrastructure protection.

4.2 Optical Internetworking Forum (OIF)

As its name indicates, the OIF is an industry forum dedicated to developing optical networking specifications. The OIF does not claim to develop “standards.” However, this has not prevented OIF from developing a number of physical layer specifications, and it is close to completing version 1.0 of the Optical User Network Interface (UNI 1.0). The UNI 1.0 specification is a major milestone in optical networking standards. The OIF has now begun to study the need for a NNI specification, but there is no consensus as to whether NNI is Node-to-Node Interface or Network-to-Network Interface.

The OIF, founded in 1998, has an interesting mixture of cultures. It has an IP-centric view, which favors IP-based specifications for network control and management protocols. However, it has high regards for the requirements of service providers, and, for this reason, views of service providers are well received. The UNI 1.0 specification represents service providers’ views, which favor bifurcation of the network into access and core, whereas the IETF does not.

In addition, the OIF coordinates better with the ITU-T and attempts to avoid duplication of efforts with ITU-T whenever possible.

In terms of the process for developing specifications, the OIF is closer to the ATM Forum than it is to the ITU-T or the IETF. It has open, contribution-driven processes and makes decisions by voting.

Opportunities

The OIF is developing a document on Operation, Administration, Management and Provisioning (OAMP) requirements. The OAM&P requirements document will have a section on security. Early acceptance of DoD requirements into the security section will be important for DoD’s efforts to produce security specifications for optical networks.

Possible near term actions by DoD include:

- A contribution to the next OIF meeting in November 2001 on the security section of the OAM&P requirements document,
- A contribution to the next OIF meeting in November 2001 on overall security requirements.

4.3 ITU-T

Since the early 1990s, the ITU-T has begun to lose its dominance over telecommunication standards. First came the ATM Forum, which quickly received recognition by industry as the place to develop ATM standards. Optical networking standards is another similar case. The IETF and the OIF receive more recognition by the

industry than the ITU-T. ITU-T Study Group 15 seems to have a dominant influence on optical transmission standards, of which G.872 and G.709 are examples.

Assessment

Because the most significant standards work by ITU-T on optical networking covers transmission standards, which have little relevance in terms of network control and management security, the ITU-T does not appear to be a primary venue for DoD's optical networking security standards work.

4.4 Possible Strategies

A reasonable strategy for promoting optical networking security standards is a two-prong approach, i.e., to work with both the IETF and the OIF at the same time. The challenge for us is how to work with both of these organizations most effectively.

IETF

Regarding the IETF, it appears that we cannot narrowly focus on the ccamp Working Group. If we had a narrow view of doing optical networking security only, and if we had a small set of security tools we could use to specify optical networking security standards, we would not have to work with groups other than ccamp. However, it appears that we need comprehensive security solutions, which span multiple working groups.

One approach would be to work with the IESG to explain our requirements and ask for advice how to accomplish the work within the IETF. This will require a presentation, which includes:

- Our requirements for infrastructure hardening,
- The need for a comprehensive security solution,
[Before approaching the IESG, we need to answer, for ourselves, why using IPsec or a certain simplified subset of IPsec to secure all of the IP-based control and management protocols is not a satisfactory solution. This is, after all, exactly what we did for ATM signaling and PNNI routing—rfg.]
- The rationale for IETF standards (as opposed to mil-std).

OIF

The OIF is much smaller and less diverse than the IETF. However, we have not convinced the OIF of the importance of this work on infrastructure hardening. This is mainly due to the industry's view that security is low priority or a simple matter that is not directly related to their bottom line.

For this reason, it may be also beneficial to approach the Board of Directors of the OIF with the same presentation described above and ask for its advice.

5. Roadmap

A near term roadmap for technical proposals to the IETF and the OIF is shown below:

- a. Framework
- b. Assessment of existing protocols for their capabilities to support our requirements
- c. Development of solutions for the deficiencies identified in step (b).

- d. Get comments on the solutions developed in step (c).
- e. Generate contributions to IETF and OIF
 - OIF by late October
 - IETF by late November
- f. Form partnerships with switch and router vendors to build consensus for a comprehensive security solution for infrastructure protection.

ANNEX A

List of Requirements

- **Authentication, data integrity, confidentiality, and replay detection shall be provided.**

The data integrity and replay detection services require expansion of the data stream, so it may not be possible to provide these services when a fixed data rate or data structure has already been specified.

- **The mechanisms that provide authentication, data integrity, replay detection, confidentiality, and any other security services shall be usable separately.**
Some applications have different security needs from others, and different political jurisdictions have varying regulations with respect to these mechanisms. The security mechanisms should be partitioned such that the functions of authentication, data integrity, replay detection, and confidentiality may be implemented and selected independently of each other, insofar as this makes sense. That is, it may not be sensible, for example, to try to provide replay detection with a counter unless the counter has integrity protection.
- **The users of these security services shall be able to negotiate which services, algorithms, and key lengths they will use. This negotiation protocol shall be protected from attacks that attempt to prevent use of the strongest common choice. [How is user defined? Is the user the person sending the information or is the user the person establishing a SLA with a network provider? I believe it is the person establishing a SLA but either way it should be defined more explicitly.]**

This requirement promotes interoperability at the highest common level of security.

- **For connection-oriented protocols using the data integrity service, the security protocol shall detect premature closure or truncation attacks.**

Data integrity means “the whole truth and nothing but the truth.”

- **Security specifications shall allow for the inclusion of new algorithms, block sizes, key lengths, and similar changes in other parameters.**

Advances in processing power and cryptanalysis require corresponding advances in protection methods.

- **A security specification should define multiple algorithms and key lengths with different tradeoffs of efficiency versus presumed strength.**
Because users, applications, and configurations have different security requirements, and certain jurisdictions limit the import, export, or use of certain cryptographic options, a range of mechanisms and algorithms should be specified.
- **Security services shall provide for negotiation of private algorithms not specifically identified in a given specification.**
Selection of algorithms is a sensitive issue for some organizations, who will want security algorithms tailored to their needs. The security mechanisms shall provide

standardized mechanisms for support of such private algorithms, including means for agreement between two users to choose a private rather than a specification-defined algorithm.

- **Two-way authentication shall be supported.**

In many protocols, authentication is tightly linked to the establishment of keys to protect subsequent traffic. Two-way authentication is often required to prevent various active attacks on the protected protocol and the secure establishment of such keying material.

- **The security option shall provide for both manual and automated key management.**

Whereas automated key management is required in most production environments, manual key management provides a backup mechanism and a method for testing security mechanisms.

- **Security specifications shall define mechanisms that will scale to large numbers of users.**

It may be appropriate to include multiple mechanisms with different scaling properties, but the choice of the mandatory-to-implement option shall take scaling into account.

- **Automated key management options shall include one based on certificates and PKI. The additional choices may include shared secret key servers or quantum key distribution.**

PKI-based systems are the ones most commonly encountered, but certain shared secret systems like Kerberos are also used.

- **A public-key-based solution should support a perfect forward secrecy option.**

Compromise of long-term keys at some future time should not compromise the confidentiality of traffic previously protected by short-term keys derived with the help of the later-compromised long-term keys.

- **Automated key management shall include options to refresh keys after a certain time period or traffic volume.**
- **Implementations of the confidentiality service should, as far as possible, protect against traffic flow analysis by hiding network addresses, higher layer protocol headers, messages lengths, and timing aspects of the protocol.**
- **Security protocols should be designed to avoid introducing and to minimize the impact of denial of service attacks.**

Some security mechanisms and algorithms require substantial processing or storage, in which case the security protocols should protect themselves as well as possible against flooding attacks that overwhelm an endpoint with such processing.

- **Security services may be integrated with the corresponding end-point protocol engine or be implemented in a separate security module.**

Both approaches have their own advantages and one may be preferred over the other in a given situation.

- **As far as practical, key management protocols and traffic protection mechanisms should conceal network addresses, routing, and other information that exposes the size, topology, or traffic patterns of the network.**
- **Nesting of security services should be supported.**

This capability allows different security services to be applied to different parts of a communications path.